


ARTIFICIAL INTELLIGENCE (AI): WHAT YOU NEED TO KNOW AND HOW IT WILL CHANGE HUMAN HISTORY

Sheldon Hochberg
Friendship Heights Village Center
October 23, 2017



“Artificial intelligence is shaping up as the next industrial revolution, poised to rapidly reinvent business, the global economy and how people work and interact with each other.”

How Artificial Intelligence Will Change Everything, Wall St. Journal, March 6, 2017

“AI is enormously disruptive and will kill jobs, but will also improve society.” *Warren Buffet, May 2017*

The possibility “of artificial intelligence taking over American jobs is so far away [that it is] not even on my radar screen.” *Steven Mnuchin, Secretary of the Treasury, March 2017*

“Artificial Intelligence is no match for natural stupidity.” *Albert Einstein, Date Unknown*

“The one who becomes the leader in this sphere will be the ruler of the world.” *Vladimir Putin, August 2017*

“I think the development of full artificial intelligence could spell the end of the human race.” *Stephen Hawking, May 2017*

The Trajectory of Technological Progress

Progress Is Not Linear.
There Are Inflection Points That
Accelerate Progress

Human Progress

Time



Three “Die” Levels of Technological Progress

- ▶ From 100,000 B.C.E. to 12,000 B.C.E. (98,000 years)
 - Development of the use of fire, language, the wheel.
- ▶ From 12,000 B.C.E. to 1900 A.D. (13,900 years)
 - Development of civilization; science and math; printing press; governments; towering churches; steam engines.
- ▶ From 1900 to 2017 (117 years)
 - World-wide use of electricity; autos; planes; telephone/radio/television; computers; the Internet; space travel; knowledge available to everyone everywhere.

Advances in Computers in 70 Years

- ▶ 1946 – ENIAC (Electronic Numeric Integrator and Calculator) – the world's first programmable computer –
could perform 20,000 multiplications per minute.
- ▶ 2016 – the Sunway TaihuLight computer in Wuxi, China – the world's most powerful computer for two years in a row –
can perform 93,000 trillion calculations per second.

Futurist Ray Kurzweil's "Law of Accelerating Returns" in Human History

- ▶ Technology accelerates at faster paces in more advanced societies than in less advanced societies.
- ▶ By 2000, our rate of advancement was five times the average rate in the 1900's.
- ▶ At this rate, another century's advancement will be achieved by 2021.
- ▶ By the 2040's, a century's worth of progress may be achieved multiple times in the same year.


In Sum:

In thinking about what the world will be like in 30 years (2047), you cannot compare it with how life was 30 years ago (1987) because technological progress is not linear.

What We Will Be Covering

- I. What is Artificial Intelligence; how it works; what it does.
- II. The history of AI and where things stand today.
- III. The promise of AI over the next decades.
- IV. The concerns that need to be addressed to ensure that AI works in the best interest of society.

The Elements of Thinking



The Elements of Thinking

- ▶ **Knowledge/Understanding:**
 - Having an ever-growing knowledge of “facts”;
 - Understanding the patterns in those facts; and, hence,
 - Understanding when things are the same and when things differ.
- ▶ **Decision-Making/Judgments/Predictions:**
 - Based on that knowledge/understanding, applying “judgment” or “reason” so as to make useful decisions – **that frequently are really predictions.**

What Is “Artificial Intelligence” (AI)

A computer program (algorithm), perhaps inside a robot, that is able to do something, or make decisions, that humans can do or make – but faster, cheaper, and with greater accuracy.

Developments That Have Produced The Beginning of The AI Golden Age

- ▶ Massive amounts of relevant/quality data available in digital form.
 - “In 2016 we produced as much data as in the entire history of humankind through 2015.” *Will Democracy Survive Big Data & Artificial Intelligence*,” Scientific American, 2017.
- ▶ Massive computing power by energy-efficient computers.
- ▶ Greater understanding of how humans think and the ability to translate that understanding into mathematics and sophisticated algorithms.

The Three Types of Artificial Intelligence

- ▶ Artificial Narrow Intelligence (“ANI”): ability to carry out a specific task (play chess; get information based on voice directions (SIRI or Alexa); spot spam email; driverless cars).
- ▶ Artificial General Intelligence (“AGI”): ability to carry out different tasks that a human could do.
- ▶ Artificial Super Intelligence (“ASI”): ability to learn from its experiences and from new data to perform a wide range of actions and to generate new computer code on its own to help achieve its objectives.

Differences Between AI Programs and “Traditional” Computer Programs?

► Inputs:

◦ Traditional Programs:

- Use letters, numbers, and symbols and limited types of communication media, such as a keyboard, mouse, or disc.

◦ AI Programs:

- Inputs to an AI program can be anything perceived by the five senses – converted to digital inputs.
- Sight – one, two, or three dimensional objects.
- Sound – spoken language, music, noise made by objects.
- Touch – temperature, smoothness, resistance to pressure. etc.
- Smell – every kind of odor.
- Taste – sweet, sour, salty, bitter foodstuffs, etc.

Differences Between AI Programs and “Traditional” Computer Programs?

► Processing:

◦ Traditional Programs:

- Manipulate the stored symbols using a set of previously defined instructions.

◦ AI Programs:

- Engage in pattern matching and problem solving, where information about the world, presented to the AI program in digital format, is used to solve complex tasks;
- Can self-learn, potentially including (down the road) developing its own new algorithms to achieve the objectives for which the program was created.

Differences Between AI Programs and “Traditional” Computer Programs?

► Output:

◦ Traditional Programs:

- Limited to alphabetical/numeric symbols communicated on a computer screen, paper, or magnetic disk.

◦ AI Programs:

- In addition to the output of traditional programs, output can be in the form of synthesized speech, visual representations, manipulation of physical objects, or movement in space.

Machine Learning

- ▶ Algorithms and processors that can **classify** and **cluster** raw input data **and that improve – learn – as they are given more data.**
 - **“Classify”**: creating or applying labels to data;
 - **“Cluster”**: identifying similarities and differences between data in the classifications.
- ▶ For example:
 - Is this email spam or not spam?
 - Does this person have cancer or not?
 - Is this a case likely to win before a jury or a case likely to lose?
 - Is this a stock likely to go up or a stock likely to go down?

Types of Machine Learning

- ▶ **Supervised Learning**: the labels for the data are programmed into the algorithm. Currently the most common form of machine learning.
- ▶ **Unsupervised Learning**: no labels are provided; the algorithm learns by itself to recognize and categorize the similarities and differences in the data.

How Does Machine Learning Work

- ▶ You create a dataset that can teach the program how to differentiate.
 - For example, you provide the algorithm with hundreds of thousands of spam email and of non-spam email (“training data”) -- so that the algorithm can detect the similarities and differences between what is spam and non-spam.
- ▶ As the program develops experience with more and more spam and non-spam emails, it sharpens its ability to see the similarities and differences, and **becomes better and better at recognizing spam in an email it has never seen before.**


“Deep Learning” or “Deep Neural Networks” – Where All the Excitement Is

The AI program automatically refines its methods, and improves its results, as it gets more data, using multiple layers of abstraction – the way the mind works.

Visualization of Brain Neurons



Deep Neural Network



www.shutterstock.com - 484275199

How Can Computers Be Programmed for Deep Learning?

“Mathematical tools such as formal logic, probability, and decision theory have yielded significant insight into the foundations of [human] reasoning and decision-making.” *Research Priorities for Robust and Beneficial Artificial Intelligence*, 2015.

“The increased computer power that is making all this possible derives . . . from the realization in the late 2000s that graphics processing units (GPUs) made by Nvidia — the powerful chips that were first designed to give gamers rich, 3D visual experiences—were 20 to 50 times more efficient than traditional central processing units (CPUs) for deep-learning computations.” Roger Perloff, *Why Deep Learning Is Suddenly Changing Your Life*, 2016.


A Deep Learning Example: What Will Your House Sell For?

- ▶ Today, a real estate agent who has sold hundreds of homes and who has experience on the thinking of buyers, gives you her best estimate, based on her experience, of what your house should sell for.

How a Deep Learning AI Program Might Work

- ▶ The program is given extensive data on the characteristics and sales prices of hundreds of thousands (or millions) of houses.
- ▶ The program is given (or develops) an initial estimate as to how the various characteristics may impact (or correlate with) the sales price.
- ▶ This initial estimate, when then applied to the database of total sales, produces estimated sales prices that are off by, for example, 15%.
- ▶ The program then runs millions of continuous slight revisions of the weights for all the factors – each revision slightly increasing the accuracy of the predictions – until they reflect the actual sales price of the houses **in the database**.
- ▶ **Tests are then run to see how the program predicts the value of future sales of houses not in the database.** If the estimates are off, mathematical and statistical procedures are available to correct the program to provide more accurate predictions.

Deep Neural Network



www.shutterstock.com - 484275199

History of AI Development

- ▶ **1921:** Czech writer Karel Čapek introduces the word "robot" in his play R.U.R. (Rossum's Universal Robots). The word "robot" comes from the Czech word "robota" (work).
- ▶ **1955:** Arthur Lee Samuel (IBM) develops **checkers-playing software program** that:
 - was the world's first self-learning software program;
 - included a "search tree" of all possible plays from any position;
 - remembered every position it had ever seen and played thousands of games against itself.

History of AI Development

- ▶ **1956:** the term “artificial intelligence” coined for a conference at Dartmouth organized by a young computer scientist John McCarthy. McCarthy develops the programming language used for AI for decades – LISP.
- ▶ **1973:** After years of promise and false starts, it was predicted that AI programs will never be more capable than a talented amateur in games.
- ▶ **1990's:** Researchers start to work on algorithms – and neural networks – that can learn the logical rules of things on their own.

History of AI Development

- ▶ 2006: [Geoffrey Hinton](#) (University of Toronto & Google) develops deep neural networks.
- ▶ 2012: [Andrew Ng](#) (Stanford & Google) begins using GPU's that enable deep neural networks to operate much faster.
- ▶ 2017: Google announces development of the TPU (Tensor Processing Unit) that is 15–30 times faster than GPU's in deep neural network operations.
- ▶ Oct. 17, 2017: Google's [AutoML](#) system has produced a series of machine-learning codes with higher rates of efficiency than codes developed by the researchers themselves.

The Importance of Games

“Games provide researchers with an effective tool for training and evaluating their AI systems. As the complexity of the games they conquer increases, so does their ability to solve real-world problems.”

-- *Games Hold the Key to Teaching Artificial Intelligence Systems: The future of AI can be mapped by the evolution of its successes in mastering game, 2017*

The Importance of Games

- *Chess* (Shannon Number: 10^{123} (atoms in the universe: 10^{80})
 - 1997 – IBM's Deep Blue (specifically developed for chess) beats Chess Master Gary Kasparov.
 - Deep Blue was capable of evaluating 100 million positions a second.
- *Jeopardy!*
 - 2011 – IBM's Watson beats Ken Jennings and Brad Rutter
 - Unlike Deep Blue, Watson was developed to deal with human language and unstructured data.

The Importance of Games

- ▶ Google's DeepMind division:
 - DeepMind Technologies founded in 2010 in England to “solve intelligence”;
 - Acquired by Google in 2014 for \$500 million;
 - Unlike IBM's Deep Blue (designed for single purpose), DeepMind uses “**reinforcement learning**” to start from scratch in self-learning, and then mastering, different games.

Google's DeepMind

- ▶ In 2015, DeepMind was loaded with 49 Atari games.
- ▶ DeepMind was provided the video pixels of the game and how the score was kept.
- ▶ After playing millions of games against itself, the system learned to play and win 29 of the 49 games, **without ever being given the rules or the objective of any of the games.**
- ▶ Considered a major advance in the development of Artificial General Intelligence (AGI).

DeepMind's AlphaGo Program

► The Ancient Chinese Game of Go

(Shannon Number: 10^{170}) (atoms in the universe: 10^{80})

Regarded as the holy grail of AI. In 2015, it was believed that it would take AI until 2025 before it could beat the best human players.

- AlphaGo was taught to play Go over several months through a combination of supervised and reinforcement learning. In supervised learning, it was shown thousands of games played by top human players.
- In May 2017 AlphaGo won a three game match against Ke Jie, who had held the world No. 1 ranking for two years. After losing, Ke Jie announced his retirement.

The Amazing AlphaGo Zero

- ▶ AlphaGo Zero started with no knowledge of Go strategy **and** no training by having seen how humans play. All it was given were the rules.
- ▶ Over three days, it played 4.9 million games against itself – getting better every hour.
- ▶ In the last day, it invented advanced strategies undiscovered by human players in the multi-millennia history of the game.
- ▶ It then played 100 games against the AlphaGo program that beat Ke Jie.
- ▶ On October 18, 2017, it was reported that AlphaGo Zero had won all 100 games.

No-Limit Holdem Poker

- ▶ In chess, Go, and other games all of the possible choices are visible on the board.
 - No Limit Holdem poker is different – and involves a much more sophisticated algorithm – because:
 - the opponents' cards are hidden;
 - the amount of the bets can range from \$1 to all in; and
 - bluffing is always present.
- ▶ Many AI experts believed the toughest test to date for AI was whether it could beat top pros in this game.
- ▶ Computer scientists at Carnegie Mellon developed an AI program called *Libratus* to compete against four top poker pros for \$200k. In January 2017, 120,000 hands were played over 20 days. *Libratus* won \$1.8 million; all four pros lost.

Growth of AI

- ▶ In 2016, interest in AI outranked all other technologies.
- ▶ According to a 2016 Infosys survey of 1,600 businesses in 7 countries;
 - 76% believed that AI would be fundamental to their future;
 - AI would contribute 39% to their annual revenues by 2020;
 - 70% believe it will result in positive changes for society,
- ▶ Many thought leaders compare AI to innovations like electricity and the Internet in terms of the change it is likely to bring.
- ▶ In 2016, sales of AI were \$644 million. By 2025, sales are estimated to be \$36 billion.

Areas of the Economy That Will Be Changed Forever

- ▶ Accounting
- ▶ Advertising
- ▶ Architecture
- ▶ Crime prevention, detection, investigation
- ▶ Cybersecurity
- ▶ Education
- ▶ Fraud Detection
- ▶ Health Care
- ▶ Investment Analysis
- ▶ Law
- ▶ Management
- ▶ Music composition
- ▶ Sales
- ▶ Shipping and Logistics
- ▶ Tax Preparation
- ▶ Teaching
- ▶ Transportation (self-driving cars/trucks; logistics)
- ▶ Warfare

Where Does AI Stand Today

- ▶ SIRI and Echo Dot (Alexa)
- ▶ Crowdsourcing navigation systems (Waze)
- ▶ Spam email blockers
- ▶ Automated response service centers
- ▶ Warnings from credit card companies about potentially fraudulent charges
- ▶ Image recognition

Natural Language Processing (NLP)

- ▶ AI programs that can understand written or spoke language in the natural and different ways humans write or speak, and respond meaningfully **in that language**.
- ▶ The key: solving the fact that many words and terms have multiple meanings or may be metaphors or puns, and that people do not speak in the same syntactical ways.
 - For example, initial mechanical translators from English to Russian interpreted “the spirit is willing, but the flesh is weak” to mean “the vodka is agreeable, but the meat is spoiled.”
- ▶ We are well on our way to instantaneous translations and programs/robots that can converse with humans in a way (almost) that humans converse with each other.

Health Care

- ▶ AI will come to dominate many areas of health care because of:
 - Global shortfall of 4.3 million doctors and nurses;
 - AI is beginning to demonstrate superiority over humans in diagnosing medical conditions and in identifying the best treatment.
- ▶ “Machine learning could be a game-changer in medicine because, unlike humans, computers don’t get tired and have an infinite capacity for learning and memorization. . . . AI can reduce the burden on doctors and nurses so they can focus on the uniquely human elements of patient care.” *Patients Are About to See a New Doctor: Artificial Intelligence*, January 2017.

Health Care Examples

- ▶ By tracking 30,000 different points on patients' hearts and 8 years of patient data, AI algorithm was able to predict which patients with pulmonary hypertension would die within a year with 80% accuracy. *London Institute of Medical Services*
- ▶ AI was able to analyze 17 different diseases with 86% accuracy on the basis of patients' breath. *American Chemical Society.*
- ▶ In 2017, using the patient's DNA and its own database of tens of millions of oncological reports and studies, IBM's Watson diagnosed a Japanese woman's rare form of cancer in 10 minutes; solving a problem that the entire hospital medical staff could not solve.

Huge Potential Benefits for Improving the Quality of Life for the Elderly

- ▶ People–friendly robot caregivers;
 - Check out “Ellie–Q” online;
- ▶ Sensors and devices in the home (or wearable) to monitor health and activity, suggest measures;
 - Monitor speech, movement, facial expression
- ▶ Intelligent walkers, wheelchairs, and exoskeletons;
- ▶ Robotic pets;
- ▶ Virtual reality headsets that let seniors “travel” to places they could not otherwise get to;
 - Tests show this reduces pain by 25%.

Transportation

- ▶ Google's autonomous vehicles and Tesla's semi-autonomous cars are driving on city streets today.
 - Google's self-driving cars have logged more than 1,500,000 miles and are completely autonomous—no human input needed.
- ▶ All car manufacturers are working on this. A recent report predicts self-driving cars to be widely adopted by 2020 (if liability issues are resolved).
- ▶ In the next 10 years we will also see self-driving and remotely controlled delivery vehicles, flying vehicles, and trucks.

Investing

- ▶ In 2001 IBM published a paper highlighting how several algorithms were able to outperform actual human stock traders.
- ▶ On October 18, 2017, the first ETF using only AI for stock selection (running on IBM's Watson platform) began trading.
- ▶ “In 2000, Goldman Sachs’ cash equities trading desk in New York employed 600 traders. Today, that operation has two equity traders, with machines doing the rest. . . . In 10 years, Goldman Sachs will be significantly smaller by head count than it is today. Expect the same to happen on every trading floor at every major financial company.”
 - -- *“Goldman Sacked: How Artificial Intelligence Will Transform Wall Street,” Newsweek, Feb. 26, 2017*

ART



Cybersecurity

- ▶ NSA's MonsterMind:
 - Project disclosed by Edward Snowden in 2014.
 - An autonomous cyberwarfare software platform that can watch international Internet connections to identify and “kill” malicious cyber attacks before they hit American infrastructure.
 - Unlike missile defense, however, MonsterMind has the ability to “fire back” at the attacker, launching a cyber counter-attack of its own.

Bringing AI to All Businesses

- ▶ Amazon, Google, and Microsoft are spending billions of dollars to be in a position to provide AI services **via the Cloud** to all businesses, including small businesses that could otherwise never afford to develop such services themselves.
- ▶ This could provide small businesses and start-ups with a competitive boost that they have never had in competing for business.

Vue.ai sells technology that analyzes pieces of clothing and automatically generates an image of the full garment on a person of any size, shape, or wearing any kind of shoes.

Only two of these models are real. The others were created by the AI program.



Are There Limits on What AI Can Do?

- ▶ Can certain human attributes be replicated by AI:
 - Intuition
 - Empathy
 - Creativity

WILL AI BRING UTOPIA OR DYSTOPIA?

**What Are The Concerns
That Have To Be Addressed**

Issues of Concern That Must Be Addressed

- ▶ Determining what society wants from AI.
- ▶ Ensuring the safety of AI programs.
- ▶ Preparing for the impact on employment and education.
- ▶ Ensuring continuous human control.
- ▶ Need for governments to develop accords to deal with major issues.

I. What Do We Want from AI?

The King Midas Dilemma

- ▶ King Midas wanted the ability to turn things into gold by touching them. However, he was not perfectly clear in his prayers about precisely what he wanted. Thus, Dionysus granted him the unwanted power to turn everything he touched into gold – his food, his son.
- ▶ Many thought leaders believe that identifying precisely what we want from AI before going much further in its development is critical for the same reason.

II. Ensuring Safety of AI Programs

- ▶ Ensuring AI programs will perform as expected?
 - No unwanted behavior or consequences.
 - No intentional manipulation by unauthorized parties (e.g., malicious software).
 - Output not affected by prejudices of the creator of the software or those who input the data used.
- ▶ How do we deal with the liability issues if something goes wrong?
- ▶ How do we ensure bad actors and countries don't use AI for evil objectives?

Safety (cont'd)

- ▶ Concerns about potentially dangerous AI programs that can (for example):
 - Create fake audio and video files;
 - Wage electronic war.

Recent Problem at Facebook

- ▶ Two AI programs, “Bob” and “Alice,” that were working together, started to “talk” to each other in sentences that the programmers did not understand.
- ▶ "I can i i everything else," Bob would say. Alice would respond with "balls have zero to me to me to me to me to me to me to me to me to.""
- ▶ After determining that the programs were using shortcuts to communicate with each other that had never been programmed, Facebook closed the programs down.

III. AI's Impact on Employment

- ▶ “50% of all jobs will be lost or replaced.” *Chief Economist, Bank of England, April 2017*
- ▶ “The equivalent of more than 1.1 billion full-time jobs, including more than 100 million in the U.S. and Europe, are associated with automatable activities.” *McKinsey Study, 2017*

But What Does History Show?

- ▶ In 1589, Queen Elizabeth, after seeing a demonstration, denied a patent for a “stocking frame knitting machine,” stating:
 - “Consider thou what the invention could do to my poor subjects. It would assuredly bring to them ruin by depriving them of employment, thus making them beggars.”
- ▶ The introduction of automobiles in daily life led to an almost total decline in horse-related jobs. However, new industries emerged resulting in an immense positive impact on employment.
 - It was not only that the automobile industry itself grew. For example, new jobs were created in the motel and fast-food industries that arose to serve motorists and truck drivers.

BUT . . .

“Unlike the Industrial Revolution and the computer revolution, the A.I. revolution is **not** taking certain jobs and replacing them with other jobs. **Instead, it is poised to bring about a wide-scale decimation of jobs — mostly lower-paying jobs, but some higher-paying ones, too.**

We are thus facing two developments that do not sit easily together: enormous wealth concentrated in relatively few hands and enormous numbers of people out of work. What is to be done?”

-- *The Real Threat of Artificial Intelligence*, Kai-Fu Lee, 2017

Need to Prepare for Impacts on Employment

- ▶ The employment question in 30 years may become: **what can humans do that AI programs or robots cannot do.**
- ▶ Need to rethink the education process and what will be needed to prepare future generations for the labor markets they will face.
- ▶ Income inequality between the very rich and the rest of society will likely increase dramatically.
- ▶ “Universal income” concepts need to be explored and discussed.

The Need for Industry/Government Cooperation

“[S]cientifically literate government planners [need to] work together with computer scientists and technologists in industry to alleviate the devastating effects of rapid technological change on the economy. The cohesion of the social order depends upon an intelligent discussion of the nature of this change, and the implementation of rational policies to maximize its general social benefit.”

— Devdatt Dubhashi and Shalom Lappin, *AI Dangers: Imagined and Real*, February 2017

Inequality In & Between Countries.

- ▶ For companies and countries, with A.I. development strength begets strength:
 - the more data you have, the better your product;
 - the better your product, the more data you can collect;
 - the more data you can collect, the more talent you can attract;
 - the more talent you can attract, the better your product.
- ▶ Some companies – and some countries – will become ultra-rich and dominant. Most other countries may become dependent on those company/countries.
- ▶ In July 2017, the Chinese government announced that it intends to be the world leader in AI by 2030.

IV. Ensuring Human Control

- ▶ As of now, AI programs – particularly deep learning, neural networks that can reach complex decisions – cannot explain how they reached their decision.
 - Will we trust the output if we don't understand how it was reached?
- ▶ “The development of full artificial intelligence could spell the end of the human race. Once humans develop artificial [super] intelligence, it would take off on its own, and redesign itself at an ever-increasing rate. Humans, who are limited by slow biological evolution, couldn't compete and would be superseded.” *Stephen Hawking* (2014)

Ethics & Governance of Artificial Intelligence Fund

- ▶ Centers at Harvard & MIT are jointly serving as “founding anchor institutions” in an effort to address the global challenges of artificial intelligence (AI) from a multidisciplinary perspective.”

Urgent Calls by Experts for Governments to Develop Policies

- ▶ Jan. 2015 open letter from Elon Musk, Steven Hawking and, subsequently, 8,000 scientists, mathematicians, and AI professionals, called for research on the potential societal impacts of AI – and possible government responses:
 - Employment and inequality;
 - Disruptions of industries;
 - Liability and “machine ethics”;
 - Preventing autonomous weapons of war;
 - Validity and safety of AI programs;
 - Security

What Should You Take Away From This Presentation?

- ▶ The importance of AI developments.
- ▶ There are good things and positive changes that will come from the growth of AI.
- ▶ There are risks that society/governments need to address.
- ▶ In a free, democratic society, all of us need to stay informed about AI developments and risks so as to have a view on the need for governmental action.